arXiv:1606.02467v1 [cs.CV] 8 Jun 2016

Point-wise mutual information-based video segmentation with high temporal consistency

Margret Keuper and Thomas Brox

Department of Computer science
University of Freiburg
Germany

Abstract. In this paper, we tackle the problem of temporally consistent boundary detection and hierarchical segmentation in videos. While finding the best high-level reasoning of region assignments in videos is the focus of much recent research, temporal consistency in boundary detection has so far only rarely been tackled. We argue that temporally consistent boundaries are a key component to temporally consistent region assignment. The proposed method is based on the point-wise mutual information (PMI) of spatio-temporal voxels. Temporal consistency is established by an evaluation of PMI-based point affinities in the spectral domain over space and time. Thus, the proposed method is independent of any optical flow computation or previously learned motion models. The proposed low-level video segmentation method outperforms the learning-based state of the art in terms of standard region metrics.

1 Introduction

Accurate video segmentation is an important step in many high-level computer vision tasks. It can provide for example window proposals for object detection [12,26] or action tubes for action recognition [14,13]. One of the key challenges in video segmentation is on handling the large amount of data. Traditionally, methods either build upon some fine-grained image segmentation [2] or supervoxel [32] method [11,10,21,22,34] or they consist in the grouping of priorly computed point trajectories (e.g. [24,19]) and transform them in a postprocessing step into dense segmentations [23]. The latter is well suited for motion segmentation applications, but has general issues with segmenting non-moving, or only slightly moving, objects. Indeed, image segmentation into small segments forms the basis for many high-level video segmentation methods like [10,22,34]. A key question when employing such preprocessing is the error it introduces. While state-of-the-art image segmentation methods [2,18,3] offer highly precise boundary localization, they usually suffer from low temporal consistency, i.e.,the superpixel shapes and sizes can change drastically from one frame to the next. This causes flickering effects in high-level segmentation methods.

In this paper, we present a low-level video segmentation method that aims at producing spatio-temporal superpixels with high temporal consistency in a

Fig. 1. Results of the proposed hierarchical video segmentation method for frame 4, 14 and 24 of the ballet sequence from VSB100 [11]. The segmentation is displayed in a *hot* color map. Note that corresponding contours have exactly the same value. Segmentations at different thresholds in this contour map are segmentations of the spatio-temporal volume.

bottom-up way. To this aim, we employ an affinity measure, that has recently been proposed for image segmentation [18]. While other, learning-based methods such as [3] slightly outperform [18] on the image segmentation task, they can hardly be transferred to video data because their boundary detection requires training data that is currently not available for videos. However, in [18], boundary probabilities are learned in an unsupervised way from local image statistics, which can be transferred to video data.

To generate hierarchical segmentations, we build upon an established method for low-level image segmentation [2] and make it applicable to video data. More specifically, we build an affinity matrix according to [18] for each entire video over space and time.

Solving for the eigenvectors and eigenvalues of the resulting affinity matrices would require an enormous amount of computational resources. Instead, we show that solving the eigensystem for small temporal windows is sufficient and even produces results superior to those computed on the full system. To generate spatio-temporal segmentations from these eigenvectors according to what has been proposed in [2] for images, we need to generate small spatio-temporal segments from the eigenvectors. We do so by extending the oriented watershed transform [2] to three dimensions. This is substantial since apart from inferring object boundaries within a single frame it also allows us to predict where the same objects are located in the next frame. We achieve this without the need to compute optical flow. Instead, temporal consistency is maintained simply by the local affinities computed between frames and the smoothness within the resulting eigenvectors.

Once we have estimated the spatio-temporal boundary probabilities, we can apply the ultrametric contour map approach from [1] on the three-dimensional data.

We show that the proposed low-level video segmentation method can compete with high-level learning-based approaches [22] on the VSB100 [11,28] video segmentation benchmark. In terms of temporal consistency, measured by the region metric VPR (volume precision recall), we outperform the state of the art.

1.1 Related work

Image Segmentation An important key to reliable image segmentation is the boundary detection. Most recent methods compute informative image boundaries with learning-based methods, either using random forests [8,17] or convolutional neural networks [4,5,30]. Provided a sufficient amount of training data, these methods improve over spectral analysis-based methods [2,18] that defined the state of the art before. However, the output of such methods provides a proxy for boundary probabilities but does not provide a segmentation into closed boundaries.
Actual segmentations can be built from these boundaries by the well-established oriented watershed and ultrametric contour map approach [1] as in [2,18,3] or, as recently proposed, by minimum cost lifted multicuts [20].
Our proposed algorithm is most closely related to the image segmentation method from [18], where the PMI-measure has been originally defined. The advantage of this measure is that is does not rely on any training data but estimates image affinities from local image statistics. We give some details of this approach in section 3.

Video Segmentation The use of supervoxels and supervoxel hierarchies has been strongly promoted in recent video segmentation methods [32,7,31,16]. These supervoxels provide small spatio-temporal segments built from basic image cues such as color and edge information. While [31] tackle the problem of finding the best supervoxel hierarchy flattening, [16] build a graph upon supervoxels to introduce higher level knowledge. Similarly [10,21,22,34] propose to build graphs upon superpixel segmentations and use learned [21,22] information or multiple high-level cues [34] to generate state-of-the-art video segmentations.
In [9], an attempt towards temporally consistent superpixels has been made on the bases of highly optimized image superpixels [2] and optical flow [6]. Similar to the proposed method, [9] try to make use of advances in image segmentation for video segmentation. However, their result still is a (temporally somewhat more consistent) frame-wise segmentation that is processed into a video segmentation by a graph-based method. In the benchmark paper [11], a baseline method for temporally consistent video segmentation has been proposed. From a state-of-the-art hierarchical image segmentation [2] computed on one video frame, the segmentation is propagated to the remaining frames by optical flow [6]. The relatively good performance of this simple approach indicates that low-level cues from the individual video frames have high potential to improve video segmentation over the current state of the art.
In [27] an extension of the method from [2] to video data has been proposed. In this work, the temporal link is established by optical flow [6] and the pixel-wise

eigensystem is solved for the whole video based on heavy gpu parallelization. Temporally consistent labelings are computed from the eigenvectors by direct spectral clustering, thus avoiding to handle the problem of temporal gradients. In contrast, our method neither needs precomputed optical flow nor does it depend on solving the full eigensystem. Instead, temporal consistency is established by an in-between-frame evaluation of the point-wise mutual information [18]. Further, we extend the oriented watershed approach from [2] to the spatio-temporal domain so we can directly follow their approach in computing the ultametric contour map [1]. In this setup, we can show that solving the eigensystem for temporal windows even improves over segmentations computed from solving the full eigensystem.

2 Method Overview

The proposed method is a video segmentation adaptation of a pipeline that has been used in several previous works on image segmentation [2,18,3] with slight variations. The key steps are given in Fig. 2. We start from the entire video sequence and compute a full affinity matrix at multiple scales. Eigenvectors are computed for these scales within small overlapping temporal windows of three frames. On the three-dimensional spatio-temporal volumes of eigenvectors, spatial and temporal boundaries can be estimated. These can be fed into the ultrametric contour map hierarchical segmentation [1] adapted for three-dimensional data.

3 Point-wise mutual information

We follow [18] in defining the point-wise mutual information (PMI) measure employed for the definition of pairwise affinities. Let the random variables A and B denote a pair of neighboring features. In [18], the joint probability $P(A, B)$ is defined as a weighted sum of the joint probability $p(A, B; d)$ of features A and B occurring with a Euclidean distance of d:

$$P(A, B) = \frac{1}{Z} \sum_{d=d_0}^{\infty} w(d) p(A, B; d). \tag{1}$$

Here,Z is a normalization constant and the weighting function w is a Gaussian normal distribution with mean-value two. The marginals of the above distribution are used to define $P(A)$ and $P(B)$. To define the affinity of two neighboring points, the direct use of this the joint probability has the disadvantage of being biased by the frequency of occurrence of A and B, i.e. if a feature occurs frequently in an image, the feature will have a relatively high probability to co-occur with any other feature. The PMI corrects for this unbalancing:

$$\mathrm{PMI}_\rho(A, B) = \log \frac{P(A, B)^\rho}{P(A)P(B)}. \tag{2}$$

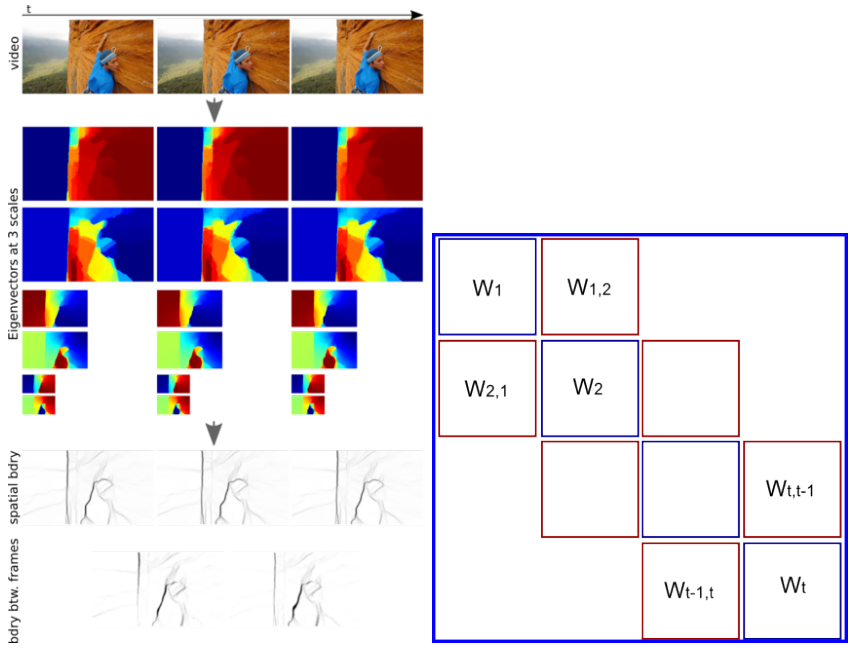

Fig. 2. Workflow of the proposed method for temporally consistent boundary detection. From an image sequence, temporally consistent eigenvectors are computed at multiple scales. While the spatial boundaries indicate object boundaries for every frame, boundaries in-between frames indicate the changes over time. In the example sequence, the *rock climbing* sequence from the VSB100 [11,28] training set, the dominant change over time in the first frames is the camera motion.

In [18], the parameter ρ is optimized on the training set of the BSD500 image segmentation benchmark [2]. We stick to their resulting parameter choice of $\rho = 1.25$.
The crucial part of the affinity measure from [18] that makes it easily applicable to unsupervised boundary detection is that $P(A, B)$ is learned specifically for every image from local image statistics. More specifically, for 10000 random sample locations per image, features A and B with mutual distance d are sampled. To model the distribution, kernel density estimation [25] is employed.

3.1 Spatio-temporal affinities

According statistics could be computed on entire videos and used to generate segmentations. However, there is a strong reason not to compute the $P(A, B)$ over all video frames: image statistics can change drastically during a video or image sequence, e.g. when new agents enter the scene, the camera moves or the illumination changes. To be robust towards these changes, we chose to estimate $P(A, B)$ per video frame and use these for the computation of affinities

within this frame within this frame and in-between this frame and the next. This is justified by the assumption that changes in local statistics are temporally smooth.
Thus, within every frame t, affinities of its elements are computed according to the estimated $\mathrm{PMI}_{\rho,t}$ for color and local variance of every pixel to every pixel within a radius of 5 pixels. Within every frame t and its successor $t+1$, affinities are computed according to $\mathrm{PMI}_{\rho,t}$ for every pixel in frame t to every pixel in frame $t+1$ within spatial distance of 3 pixels, and, for every pixel in frame $t+1$ to every pixel in frame t within the same distance. The result is a sparse symmetric spatio-temporal affinity matrix W as given in Fig. 2.

4 Spectral Boundary Detection

Given an affinity matrix W, spectral clustering can be employed to generate boundary probabilities [2,18] and segmentations [9,10,22] according to a balancing criterion, more precisely, approximating the normalized cut

$$\mathrm{NCut}(A,B) = \frac{\mathrm{cut}(A,B)}{\mathrm{vol}(A)} + \frac{\mathrm{cut}(A,B)}{\mathrm{vol}(B)}, \tag{3}$$

with $\mathrm{cut}(A,B) = \sum_{i\in A, j\in B} w_{ij}$ and $\mathrm{vol}(A) = \sum_{i\in A, j\in \mathcal{V}} w_{ij}$.
Approximate solutions to the normalized cut are induced by the first k eigenvectors of the normalized graph Laplacian $L_{\mathrm{sym}} = I - D^{-\frac{1}{2}} W D^{-\frac{1}{2}}$, where D is the diagonal degree matrix of W computed by $d_{ii} = \sum_{j\in V} w_{ij}$.
However, the computation of eigenvectors for large affinity matrices rapidly becomes expensive both in terms of computation time and memory consumption. To keep the computation tractable, we can reduce the computation to small temporal windows and employ the spectral graph reduction [10] technique.

Spectral Graph Reduction Spectral graph reduction [10] is a means of solving a spectral clustering or normalized cut problem on a reduced set of points. In this setup, the matrix W defines the edge weights in a graph $G = (V, E)$. Given some pre-grouping of vertices by for example superpixels or must-link constraints, [10] specify how to set these weights in a new graph $G' = (V', E')$ where V' represents the set of vertex groups and E' the set of edges in between them such that the normalized cut objective does not change. They show on low-level image segmentation as well as on high-level video segmentation the advantages of this method.
Since we want to remain as close to the low-level problem as possible, we employ a setup similar to the one proposed in [10] for image segmentation. More specifically, we compute superpixel at the finest level produced by [3] for every frame, which builds upon learned boundary probabilities from [8]. In order not to lose accuracy in boundary localization, [10] proposed to keep single pixels in all regions with high gradients and investigate the trade-off between pixels and

Fig. 3. Examples of boundary estimates (center) and motion boundary estimates (right) on three sequences of VSB100 [11,28] (*arctic kayak*, *riverboat*, and *salsa*). The motion boundaries can be directly derived from the proposed method.

superpixels that is necessary. Similarly, we keep single pixels in all regions with high boundary probability.

Multiscale Approach and Boundary Detection Since it has been shown in the past that spectral clustering based methods benefit from multi-scale information, we build affinity matrices also for videos spatially downsampled by factor 2 and factor 4. In this case, no pixel pre-grouping is necessary. For all three scales, we solve the eigensystems individually and compute the smallest 20 eigenvalues and according eigenvectors (compare Fig. 2). Note that these eigenvectors are highly consistent over the temporal dimension. We upsample these vectors to the highest resolution and compute oriented edges in 9 directions, with the standard oriented edge filters in 8 sampled spatial orientations and only one temporal gradient, i.e. there is no mixed spatio-temporal gradient. Depending on the frame-rate, using finer orientation sampling would certainly make sense. However, on the VSB100 dataset [11,28], we found that this simple setup works best. Examples of our extracted boundary estimates are given in Fig. 3. Visually, the estimated boundaries look reasonable and are temporally highly consistent. They form the key to the final hierarchical segmentations.

Fig. 4. Resulting UCMs from the proposed method on the sequences *arctic kayak* and *riverboat* of the VSB100 [11,28] dataset. We show the UCM value in frames 1, 11, and 21 while these actually extend over the full lengths of the sequences. Especially in the high levels of the hierarchy, the segmentations are very consistent.

Evaluation of Temporal Boundaries On the BVSD [28] dataset, benchmark annotations for occlusion boundaries were provided. This data can be used as a proxy to evaluate our temporal boundaries. Occlusion boundaries are object boundaries that occlude other parts of the scene - as opposed to within-object boundaries. In [28], the importance of motion cues for such occlusion boundaries has been pointed out. In fact, our temporal boundaries indicate boundaries separating regions within one frame that will undergo occlusion or disocclusion between this frame and the next. Thus, they can only provide part of the necessary information for object boundary detection. To extract this motion cue from our data, we apply a pointwise multiplication of the spatial boundaries at frame t with the temporal boundaries between frame t and its two neigh boring frames. Thus, if an object does not move in one of the frames, the respective edges are removed. Examples of the resulting motion boundary estimates are given in Fig. 3.

When we evaluate on the closed boundary annotations of this benchmark with the benchmark parameters from BSDS500 [2], we get a surprisingly low f-measure score of 0.34 with the best common threshold for the whole dataset, 0.41 if we allow individual thresholds per sequence. The reason might be the relatively low spatial localization accuracy of our boundaries. We ran all our experiments on the *half resolution* version of the VSB100 benchmark such that, to evaluate on the annotations from [28], boundary estimates need to be upsampled.

5 Closed Spatio-temporal Contours

Given spatio-temporal boundary estimates, closed regions could be generated by different methods such as region growing, agglomerative clustering [15], watersheds [29] or the recently proposed minimum cost lifted multicuts [20]. In [1] a mathematically sound and widely used (e.g. in [2,27,9,18,3]) setup for the generation of hierarchical segmentations from boundary probabilities and an initial fine-grained segmentation is given. The therein defined Ultrametric Contour Map provides for a duality between the saliency of a contour and the scale of its disappearance from the hierarchy.
The approach from [1] can directly be applied to three-dimensional data. The difference is that the region contours are now two-dimensional curves that meet each other in one-dimensional curves or points. Each one-dimensional curve is common to at least three contours. As in the two-dimensional case, every contour is separating exactly two regions.
Samples from the resulting Ultrametric Contour Maps can be seen in Fig. 4. The brightness of the contour, displayed in *hot* color maps, indicate the saliency of a contour, i.e. its hierarchical level in the segmentation. Over all frames of the videos, the resulting closed contours have consistent saliency.

6 Experiments and Results

Setup We compute PMI-based affinity matrices W_s on color and local variance within all frames and between every frame and its successive frame as described in section 3.1 for three different scales (1, 0.5 and 0.25). For scale 1, we employ spectral graph reduction [10], reducing the number of nodes by factor 12-15. At each scale s, we solve the eigenvalue problems of the normalized graph Laplacians corresponding to W_s for overlapping temporal windows with stride 1 to generate the first 20 eigenvectors. The best choice of the temporal window size is not obvious because of the eigenvector leakage problem, also mentioned in [27]. In the spectrally reduced graph, spatial leakage is probably low [10], so we we hope for an accordingly low temporal leakage and choose a larger temporal window of size 5, while we solve the eigenvalue problem for smaller temporal windows of length 3 for scales 0.5 and 0.25. The resulting eigenvectors are resampled to the original resolution. The average oriented gradients on these eigenvectors for the multiscale boundary estimates we use. We compare the 3D ultrametric contour maps computed from these boundary estimates to those computed on only the original scale with spectral graph reduction. For the original scale, we also compare to the results we get by solving the eigensystem on the full video without temporal windows.

Evaluation We evaluate the proposed video segmentation method on the *half resolution* version of the VSB100 video segmentation benchmark [11,28]. It consists of 40 train and 60 test sequences with a maximum length of 121 frames. Human segmentation annotations are given for every 20th frame. Two evaluation metrics are relevant in this benchmark, denoted as boundary precision and

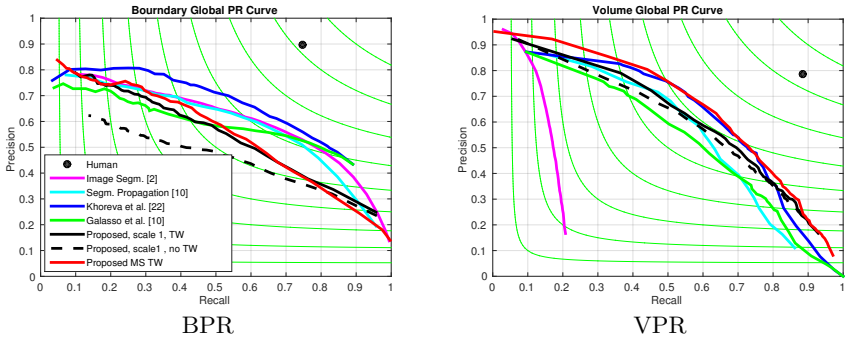

Fig. 5. Results on the general benchmark of the VSB100 [11,28] dataset. While the proposed method performs worse than the state of the art in terms of boundary precision and recall (BPR), we outperform all competing methods on the region metric VPR.

VSB100: general benchmark

	BPR			VPR		
Algorithm	ODS	OSS	AP	ODS	OSS	AP
Human	0.71	0.71	0.53	0.83	0.83	0.70
*Galasso et al. [9]	0.52	0.56	0.44	0.45	0.51	0.42
*Grundmann et al. [16]	0.47	0.54	0.42	0.52	0.55	0.52
*Ochs and Brox [23]	0.14	0.14	0.04	0.25	0.25	0.12
Xu et al. [33]	0.40	0.48	0.33	0.45	0.48	0.44
* Galasso et al. [10]	0.62	0.65	0.50	0.55	0.59	0.55
* Khoreva et al. [22] SC	**0.64**	**0.70**	**0.61**	0.63	**0.66**	0.63
Segmentation Propagation [11]	0.60	0.64	0.57	0.59	0.62	0.56
IS - Arbelaez et al. [2]	0.61	0.65	0.61	0.26	0.27	0.16
Oracle & IS - Arbelaez et al. [2]	0.61	0.67	0.61	0.65	0.67	0.68
* Proposed MS TW	0.56	0.63	0.56	**0.64**	**0.66**	**0.67**

Table 1. Comparison of state-of-the-art video segmentation algorithms [16,23,9,33,10,22,11] with the proposed low-level method. It shows the boundary precision-recall (BPR) and volume precision-recall (VPR) at optimal dataset scale (ODS) and optimal segmentation scale (OSS), as well as the average precision (AP). The algorithms marked with (*) have been evaluated on resized video frames to 0.5 of the original resolution.

recall (BPR) and the volume precision recall (VPR). The BPR measures the accuracy of the boundary localizations per frame. Image segmentation methods usually perform well on this measure, since temporal consistency is not taken into account. The VPR is a region metric. Here, exact boundary localization is less important, while the focus lies on the temporal consistency. This is the measure on which we expect to perform well.

VSB100: motion subtask

	BPR			VPR		
Algorithm	ODS	OSS	AP	ODS	OSS	AP
Human	0.63	0.63	0.44	0.76	0.76	0.59
*Galasso et al. [9]	0.34	0.43	0.23	0.42	0.46	0.36
*Ochs and Brox [23]	0.26	0.26	0.08	0.41	0.41	0.23
* Khoreva et al. [22]	0.45	**0.53**	0.33	0.56	**0.63**	0.56
Segmentation Propagation [11]	**0.47**	0.52	**0.34**	0.52	0.57	0.47
IS - Arbelaez et al. [2]	0.47	0.43	0.35	0.22	0.22	0.13
Oracle & IS - Arbelaez et al. [2]	0.47	0.34	0.35	0.59	0.60	0.60
* Proposed MS TW	0.41	0.43	0.29	**0.58**	0.58	**0.58**

Table 2. Comparison of state-of-the-art video segmentation algorithms with the proposed low-level method on the motion subtask of VSB100 [11].

Results The results of our PMI-based video segmentation are given in Fig. 5 in terms of BRP and VPR curves for 51 different levels of segmentation granularity, which is the standard for the VSB100 video segmentation benchmark. In terms of BPR, all our results remain below the state of the art. There can be several reasons for this behavior: (1) Competitive methods that are based on per image superpixels usually compute these superpixels on the highest possible resolution [10,22], while we start from the *half resolution* version of the benchmark data. (2) The number of eigenvectors we compute might be too low. In our current setup, we compute the first 20 eigenvectors per matrix. The optimal number to be chosen here varies strongly depending on the structure of the data as well as the employed affinities. (3) Most importantly, by the definition of our boundary detection method, we force these boundaries to be spatially consistent. However, spatial consistency is not at all required for this measure. If our boundaries show some amount of temporal smoothness, this might cause the boundaries to be slightly shifted in an individual frame.

On the VPR, the proposed method benefits from the temporal consistency it optimizes and outperforms all previous methods in the three aggregate measures ODS (meaning, we choose one global segmentation threshold for the whole dataset), in OSS, which allows to choose the best threshold per sequence, and average precision (AP). Respective numbers are given in Tab. 1. This gives rise to the conclusion that the proposed segmentations are indeed temporally consistent.

In Fig. 5, we also plot the result we get if we only use boundary estimates from the original resolution without using multiscale information (depicted in black). As we expect, the segmentation quality remains below the quality we get with the multiscale approach. However, we did not know a priori what to expect from solving the eigensystem for the entire videos without applying temporal windows (dashed black line). To do so actually requires large amounts of memory (> 512Gb) for most sequences. The results actually remain clearly below those computed with temporal windows. In fact, the eigenvectors we compute on the

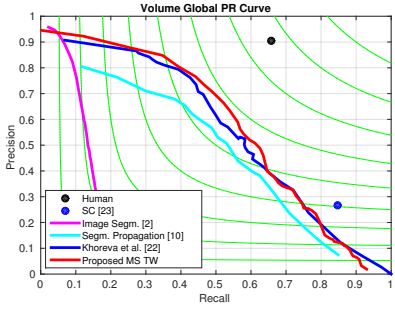

Fig. 6. Results on the motion subtask the VSB100 [11,28] dataset in terms of the region metric VPR.

small temporal windows show high temporal consistency while those computed on the whole video are subject to temporal leakage of the eigenvectors, meaning that values in the eigenvectors within a region can change smoothly throughout the sequence, resulting in decreased discrimination power.
The best results in Tab. 1 are those from the oracle case, where the individual image segments from [2] are temporally linked based on the ground truth. The numbers indicate the best result one could achieve on the benchmark, starting from the given image segmentation.

Results on Motion Segmentation Since our segmentations claim to be temporally consistent even under significant motion in videos, we also performed the evaluation on the motion subtask of VSB100. For this motion subtask, only a subset of the videos, showing significant motion, is evaluated. Non-moving objects within these videos are not taken into account. Results are reported in Tab. 2 and Fig. 6. As for the general benchmark, our results are outperformed by the state of the art on the BPR but improve over the state of the art in terms of temporal consistency, measured by the region metric VPR.

7 Conclusions

We have proposed a method for computing temporally consistent boundaries in videos. To this end, the method builds spatio-temporal affinities based on pointwise mutual information at multiple scales. On the video segmentation benchmark VSB100, the resulting hierarchy of spatio-temporal regions outperforms state-of-the-art methods in terms of temporal consistency as measured by the region metric VPR. We believe that the coarser hierarchy level can help extract high-level content of video. The finer hierarchy levels can serve as temporally consistent spatio-temporal superpixels for learning based video segmentation or action recognition.

Acknowledgments

We acknowledge funding by the ERC Starting Grant VideoLearn.

References

1. P. Arbelaez. Boundary extraction in natural images using ultrametric contour maps. In *CVPR workshop*, 2006.
2. P. Arbeláez, M. Maire, C. C. Fowlkes, and J. Malik. Contour detection and hierarchical image segmentation. *IEEE TPAMI*, 33(5):898–916, 2011.
3. P. Arbeláez, J. Pont-Tuset, J. Barron, F. Marques, and J. Malik. Multiscale combinatorial grouping. In *CVPR*, 2014.
4. G. Bertasius, J. Shi, and L. Torresani. Deepedge: A multi-scale bifurcated deep network for top-down contour detection. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
5. G. Bertasius, J. Shi, and L. Torresani. High-for-low and low-for-high: Efficient boundary detection from deep object features and its applications to high-level vision. *CoRR*, abs/1504.06201, 2015.
6. T. Brox and J. Malik. Large displacement optical flow: descriptor matching in variational motion estimation. *IEEE TPAMI*, 33(3):500–513, 2011.
7. J. Chang, D. Wei, and J. W. F. III. A Video Representation Using Temporal Superpixels. In *IEEE Computer Vision and Pattern Recognition Conference on Computer Vision*, 2013.
8. P. Dollár and C. L. Zitnick. Structured forests for fast edge detection. In *ICCV*, 2013.
9. F. Galasso, R. Cipolla, and B. Schiele. Video segmentation with superpixels. In *ACCV*, 2012.
10. F. Galasso, M. Keuper, T. Brox, and B. Schiele. Spectral graph reduction for efficient image and streaming video segmentation. In *CVPR*, 2014.
11. F. Galasso, N. Nagaraja, T. Cardenas, T. Brox, and B.Schiele. A unified video segmentation benchmark: Annotation, metrics and analysis. In *ICCV*, 2013.
12. R. Girshick. Fast R-CNN. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2015.
13. G. Gkioxari, R. Girshick, and J. Malik. Actions and attributes from wholes and parts. In *ICCV*, 2015.
14. G. Gkioxari and J. Malik. Finding action tubes. In *CVPR*, 2015.
15. R. C. Gonzalez and R. Woods. *Digital Image Processing 2nd Edition*. 2002.
16. M. Grundmann, V. Kwatra, M. Han, and I. Essa. Efficient hierarchical graph based video segmentation. In *IEEE CVPR*, 2010.
17. S. Hallman and C. C. Fowlkes. Oriented edge forests for boundary detection. In *CVPR*, 2015.
18. P. Isola, D. Zoran, D. Krishnan, and E. H. Adelson. Crisp boundary detection using pointwise mutual information. In *ECCV*, 2014.
19. M. Keuper, B. Andres, and T. Brox. Motion trajectory segmentation via minimum cost multicuts. In *ICCV*, 2015.
20. M. Keuper, E. Levinkov, N. Bonneel, G. Lavoue, T. Brox, and B. Andres. Efficient decomposition of image and mesh graphs by lifted multicuts. In *ICCV*, 2015.
21. A. Khoreva, F. Galasso, M. Hein, and B. Schiele. Learning must-link constraints for video segmentation based on spectral clustering. In *GCPR*, 2014.
22. A. Khoreva, F. Galasso, M. Hein, and B. Schiele. Classifier based graph construction for video segmentation. In *CVPR*, 2015.

23. P. Ochs and T. Brox. Object segmentation in video: a hierarchical variational approach for turning point trajectories into dense regions. In *ICCV*, 2011.
24. P. Ochs, J. Malik, and T. Brox. Segmentation of moving objects by long term video analysis. *IEEE TPAMI*, 36(6):1187 – 1200, Jun 2014.
25. E. Parzen. On estimation of a probability density function and mode. *Ann. Math. Statist.*, 33(3):1065–1076, 09 1962.
26. S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NIPS*, 2015.
27. N. Sundaram and K. Keutzer. Long term video segmentation through pixel level spectral clustering on gpus. In *ICCV Workshops*, 2011.
28. P. Sundberg, T. Brox, M. Maire, P. Arbelaez, and J. Malik. Occlusion boundary detection and figure/ground assignment from optical flow. In *CVPR*, 2011.
29. L. Vincent and P. Soille. Watersheds in digital spaces: an efficient algorithm based on immersion simulations. *IEEE TPAMI*, 13:583–598, 1991.
30. S. Xie and Z. Tu. Holistically-nested edge detection. *CoRR*, abs/1504.06375, 2015.
31. C. Xu and J. Corso. Evaluation of super-voxel methods for early video processing. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2012.
32. C. Xu, S. Whitt, and J. Corso. Flattening supervoxel hierarchies by the uniform entropy slice. In *Proceedings of the IEEE International Conference on Computer Vision*, 2013.
33. C. Xu, C. Xiong, and J. Corso. Streaming hierarchical video segmentation. In *Proceedings of European Conference on Computer Vision*, 2012.
34. S. Yi and V. Pavlovic. Multi-cue structure preserving MRF for unconstrained video segmentation. *CoRR*, abs/1506.09124, 2015.